

SYSTEM AND METHOD FOR RUNTIME  
REALLOCATION OF PLD RESOURCES  
Neil G. Jacobson

FIELD OF THE INVENTION

The present invention generally relates to runtime reconfigurable systems, and more particularly to runtime reallocation of PLD resources in response to changing activity levels.

BACKGROUND

The field of reconfigurable computing has advanced steadily for the past decade, using programmable logic devices (PLDs) as the basis for high-performance reconfigurable systems. Run-Time Reconfigurable (RTR) systems distinguish themselves by performing circuit logic customization at run-time. RTR systems using PLDs are expected to result in systems that require less hardware, less software, and fewer input/output resources than traditional ASIC or PLD-based systems.

In multiprocessor data processing systems, various systems are known for balancing the work load between the processors. An advantage to such systems is that the functions implemented by the system can be software controlled, thereby providing great flexibility. However, such systems tend to be large, expensive, and complex. Furthermore, software-implemented functions tend to be slower than hardware implementations.

A system and method that address the aforementioned problems, as well as other related problems, are therefore desirable.

EL539650972US)

1 SUMMARY OF THE INVENTION

2 A system and method for allocating resources of  
3 programmable logic devices (PLDs) according to the activity  
4 levels of the functions implemented on the PLDs are  
5 provided. A load monitor surveys the activity levels of the  
6 functions implemented on the PLDs. When decreasing and/or  
7 increasing activity levels are detected, the PLD resources  
8 are reallocated between the various functions in proportion  
9 to the decreasing and/or increasing activity levels.

10 In another embodiment, a reserve of PLD resources is  
11 maintained. When the activity level of one or more  
12 functions is increasing and none of the other functions have  
13 decreasing activity levels, some portion of the reserve is  
14 allocated to the functions having increasing activity  
15 levels.

16 Various other embodiments are set forth in the Detailed  
17 Description and Claims which follow.

18  
19 BRIEF DESCRIPTION OF THE DRAWINGS

20 FIG. 1 is a functional block diagram of a system 100  
21 having control logic for reconfiguration of PLD resources  
22 and load monitoring of different functions that are  
23 implemented on one or more PLDs, in accordance with one  
24 embodiment of the invention.

25 FIG. 2 is a flowchart of a process for monitoring the  
26 activity level of the functions implemented on the PLD(s).

27 FIG. 3A is a flowchart of a process for determining  
28 the levels at which PLD resources should be reallocated  
29 between a number of functions when it is detected that the  
30 activity level for one of the functions is decreasing.

31 FIG. 3B is a flowchart of a process for determining  
32 the levels at which PLD resources should be allocated from  
33 the reserve of resources when the activity level of a  
34 function is increasing.

FIG. 4 is a flowchart of a process for controlling reconfiguration of PLDs in response to control parameters from the load monitor.

#### DETAILED DESCRIPTION

In various embodiments, the invention monitors the activity level of various functions that are implemented on one or more PLDs. When the activity level of a function has decreased, a determination is made as to whether PLD resources of the function should be reprogrammed for another function. If so, the PLD resources are reprogrammed accordingly.

FIG. 1 is a functional block diagram of a system 100 having control logic for reconfiguration of PLD resources and load monitoring of different functions that are implemented on one or more PLDs, in accordance with one embodiment of the invention. The system includes one or more PLDs 102 that implement different functions, a storage element 104 for storing one or more configuration bitstreams 106 that program the PLD(s) to implement the functions, a configuration control element 108, and a load monitor 110.

The invention is applicable to systems having any of a variety of PLDs, for example FPGAs or CPLDs, where the particular function and the resources available on the PLD will determine the number of PLDs required. Partial reconfigurability of the PLD will also influence the number PLDs in the system. For example, the Virtex FPGA from Xilinx supports partial reconfiguration. Therefore, multiple functions may be implemented on a Virtex FPGA, depending on the resource requirements of the different functions. If it is determined that FPGA resources should

1 be taken from one function and devoted to another function,  
2 the FPGA can be partially reconfigured for the reallocation  
3 of resources.

4 The different functions are implemented with  
5 associated configuration bitstreams. For example,  
6 configuration bitstream 106 implements a first function,  
7 configuration bitstream 112 implements a second function,  
8 configuration bitstream 114 implements a third function,  
9 and configuration bitstream 116 implements an *n*th function.

10 In one embodiment, storage 104 for the configuration  
11 bitstreams may include one or more in-system memory  
12 elements. In another embodiment, the configuration  
13 bitstreams may be stored outside the system on network  
14 storage, for example. Those skilled in the art will  
15 recognize other suitable storage alternatives for different  
16 application requirements.

17 Load monitor 110 is coupled to the PLDs for monitoring  
18 the activity level of the functions. For example, the PLDs  
19 may implement a plurality of communications channels,  
20 wherein different ones of the channels are dedicated to  
21 different protocols. If the traffic on the channels  
22 implementing a first protocol has decreased and the traffic  
23 on the channels implementing a second protocol has  
24 increased, some of the channels for the first protocol may  
25 be reallocated for the second protocol. The load monitor  
26 not only detects the need for reprogramming but also  
27 determines which PLD resources should be reprogrammed.  
28 Load monitor 110 also signals the functions implemented on  
29 the PLDs when a reconfiguration is to occur so that work-  
30 in-process can be finished or reassigned.

31 The mechanism by which load monitor 110 monitors the  
32 activity levels of the different functions can be adapted

1 in accordance with the requirements of the application.  
2 For example, the functions may be designed with a feature  
3 for reporting the activity levels via an output port of the  
4 PLDs. Alternatively, registers internal to the PLDs and  
5 readable by the load monitor may be used for storing the  
6 current activity levels.

7 In order to reconfigure PLD resources, load monitor  
8 110 must track which functions are implemented on which PLD  
9 resources. Furthermore, if the different functions require  
10 different types and numbers of PLD resources, those too  
11 must be tracked by the load monitor.

12 Configuration control element 108 is coupled to load  
13 monitor 110, storage 104, and to the PLDs 102. When  
14 reprogramming is requested by the load monitor,  
15 configuration control 108 signals the selected PLD(s) and  
16 initiates transfer of the selected configuration  
17 bitstream(s) from storage 104 to the appropriate PLD(s).  
18 It will be appreciated that the logic of load monitor 110  
19 and configuration control element 108 could be integrated  
20 into a single device in another embodiment. The particular  
21 distribution of functionality between load monitor 110 and  
22 configuration control element 108, as well as the  
23 particular communication links to PLDs 102, can be tailored  
24 for specific application requirements.

25  
26 FIG. 2 is a flowchart of a process for monitoring the  
27 activity level of the functions implemented on the PLD(s).  
28 The process generally entails performing an initial  
29 allocation of PLD resources between the functions and an  
30 initial configuration of the PLDs and then monitoring  
31 activity levels of the functions and reconfiguring the PLDs  
32 as necessary.

1       At step 152, an initial allocation of PLD resources is  
2       selected for the various functions. For example, in an  
3       application such as a communications switch, certain  
4       numbers of channels may be dedicated to different  
5       communications protocols. Thirty-two channels may be  
6       dedicated to protocol A, 64 channels may be dedicated to  
7       each of protocols B and C, and so on. Thus, the load  
8       monitor determines the number of channels to devote to each  
9       of the protocols.

10       Each channel maps to a set of PLD resources. In one  
11       embodiment, not all resources are initially allocated. For  
12       example, 10% of the total channels are not allocated to any  
13       protocol and are held as a reserve of resources for peak  
14       activity levels. At step 154, the configuration control is  
15       signaled to begin the initial configuration. The load  
16       monitor provides the configuration control with information  
17       identifying which PLD resources are to be configured with  
18       which configuration bitstreams. The manner in which PLD  
19       resources and configuration bitstreams are identified will  
20       depend on the specific functions and the capabilities of  
21       the particular PLD. Furthermore, if a single PLD supports  
22       more than one function, then a number of configuration  
23       bitstreams may be required to cover the possible  
24       combinations for allocating resources of the PLD between  
25       the different functions.

26       Once the PLDs have been configured and the system is  
27       operational, at step 156, the load monitor establishes  
28       communications with the functions on the PLDs. At step  
29       158, the load monitor begins monitoring the activity levels  
30       of the functions. The load monitor samples (step 160) the  
31       activity levels every  $t$  interval of time for each function.  
32       A typical measure of activity level for a communications

1 protocol channel is the number of packets processed per  
2 second. If a protocol is supported across multiple  
3 channels, the activity level is the total activity across  
4 all the channels implementing the protocol.

5       Decision step 162 tests whether the activity level of  
6 any of the functions is decreasing. When a function shows  
7 a decreasing activity level over  $n$  samples and the overall  
8 decrease is greater than a given threshold, some of the  
9 resources allocated to the function are freed for re-  
10 allocation to other applications. The value of  $n$  and the  
11 threshold level are chosen so that reallocation is not  
12 performed in response to random and instantaneous changes.  
13 Each function may have a different threshold that varies  
14 based on the number of resources allocated to the function  
15 since functions having greater numbers of resources will be  
16 able to accommodate greater activity levels. In various  
17 embodiments, the functions implemented on the PLDs are  
18 programmed to periodically report activity levels to the  
19 load monitor either via an on-board PLD register or by  
20 transmission to the load monitor.

21       Step 164 determines how resources are to be  
22 reallocated between the functions implemented on the PLDs  
23 (FIG. 3A describes how the reallocated resource levels are  
24 determined for decreasing activity levels). At step 166,  
25 the load monitor signals the PLD(s) that will be  
26 reconfigured to complete work-in-process prior to  
27 reconfiguration. The work may be completed by virtue of  
28 finishing the work or transferring the work to a different  
29 resource (e.g., another PLD that will not be reconfigured).  
30 For I/O tasks on resources to be freed, a number of steps  
31 are performed:

1        1. The output ports from the resources to be freed are  
2        connected to output ports of another resource that  
3        implements the same function. For example, if  
4        channels A and B both implement protocol 1 and the  
5        resources from channel A are to be freed, the output  
6        ports of channel A are connected to unassigned  
7        output ports of channel B, or if time-slicing is  
8        used, to outputs of channel B having unused  
9        capacity.

10       2. Direct input data from the resource to be freed to  
11       another resource that implements the same function.  
12       For example, (continuing the example of step 1) the  
13       input data directed at channel A is redirected to  
14       the input ports of channel B.

15       3. Disable the output ports of the resource to be  
16       freed.

17       4. Disable the input ports of the resource to be freed.

18  
19       The connection and redirection of inputs and outputs can  
20       best be facilitated by either special I/O circuitry (for  
21       example, a crossbar switch), or a PLD reconfigured to re-  
22       map input and output connections.

23       At step 168, the configuration control element is  
24       signaled to begin the reconfiguration. The PLD resources  
25       that are reconfigured are those identified at step 164.

26       If decision step 162 finds that none of the functions  
27       have decreasing activity levels, the process is directed to  
28       decision step 172 to test whether any of the functions have  
29       increasing activity levels. If there are no functions  
30       having decreasing or increasing activity levels, the  
31       process returns to sampling the activity levels at step  
32       160.



1        If any of the functions are detected as having  
2        increased activity levels, the process is directed to step  
3        174. When a function shows an increasing activity level  
4        over  $n$  samples and the overall increase is greater than a  
5        given threshold, resources from the reserve are allocated to  
6        the function(s). The values of  $n$  and the threshold level  
7        are chosen so that reallocation is not performed in  
8        response to random and instantaneous changes. The  
9        thresholds associated with increasing activity levels may  
10       be different from the thresholds associated with decreasing  
11       activity levels, and each function may have a different  
12       threshold that varies based on the number of resources  
13       allocated to the function since functions having greater  
14       numbers of resources will be able to accommodate greater  
15       activity levels.

16       At step 174, it is determined how resources from the  
17       reserve are to be allocated to the function(s) (FIG. 3B  
18       describes how the reallocated resource levels are  
19       determined for increasing activity levels). The process  
20       then continues at step 168 as described above.

22       FIG. 3A is a flowchart of a process for determining  
23       the levels at which PLD resources should be reallocated  
24       between a number of functions when the activity level of a  
25       function is decreasing. The number of resources to free  
26       from a function is determined at step 202 and is based on  
27       the following model. Each resource can service a maximum  
28       load of  $R_L$  for its given function. Thus, if  $M$  resources are  
29       assigned to a function, then the maximum sustainable  
30       activity level is  $M \cdot R_L$ . If the load monitor determines that  
31       the activity level has decreased such that a current  
32       activity level =  $(M-P) \cdot R_L$  ( $P \leq M$ ), then the  $(P-1)$  resources

1 are freed (one is left in reserve with the function). At  
2 step 204, the number of resources of the function is  
3 reduced by the amount determined at step 202. For example,  
4 the number of channels that implement a particular protocol  
5 is reduced.

6 Decision step 206 determines whether the freed  
7 resources should be allocated to other functions or held in  
8 reserve. The decision is made based on a projection of  
9 future activity levels of the other functions. The  
10 projection can be linear, least squares, or another  
11 technique that best models the overall system behavior.  
12 In one embodiment, the projection is made based on past  
13 behavior. Sample activity levels are taken by the load  
14 monitor. The values associated with the activity levels  
15 are stored over a sufficiently long window of time. The  
16 length of the window is based on the usage environment.  
17 Past performance is modeled statistically to predict future  
18 behavior. If the predicted future behavior is either a  
19 decrease or increase in the activity level, then  
20 reconfiguration and re-assignment may be initiated.

21 At step 208, the allocation of freed resources to  
22 functions is determined based on the projected,  
23 proportional needs of the functions. For example, if the  
24 activity level of a first function is projected to increase  
25 by 30% and the activity level of a second function is  
26 projected to increase by 10%, then the ratio of the  
27 projected increases is used when making the allocation.  
28 Assuming both the first and second functions begin with the  
29 same initial resource allocation, then the first function  
30 will be allocated three-fourths of the resources and the  
31 second function is allocated one-fourth of the resources.

1 If none of the activity levels is increasing, decision  
2 step 206 directs control to step 210, and the freed  
3 resources are added to the reserve of free resources.  
4 Parameters that describe the resources to allocate to the  
5 different functions are then returned to the process of  
6 FIG. 2.

7  
8 FIG. 3B is a flowchart of a process for determining  
9 the levels at which PLD resources should be allocated from  
10 the reserve when the activity level of a function is  
11 increasing. At step 222, the number of resources to remove  
12 from the reserve and allocate to the one or more functions  
13 is determined.

14 The allocation of resources takes into account the  
15 increasing needs of the entire system (i.e., across  
16 multiple functions). The allocation process also ensures  
17 that the entire reserve is not depleted so that future  
18 increased activity levels can be handled. In one  
19 embodiment, the following model is used to manage the  
20 reserve of resources:

- 21 1. The load monitor tracks the total percentage of  
22 resources allocated to each function and the  
23 percentage of resources presently in the reserve. The  
24 total of all allocated resources is 100%  
25  $(R_1 + R_2 + R_3 + R_4 + \dots + R_{\text{RESERVE}} = 100\%, \text{ where } R_x \text{ is the percent}$   
26  $\text{resources allocated to function } X).$
- 27 2. The load monitor tracks the predicted increase (or  
28 decrease) of resources for each function. Each  
29 predicted increase (or decrease) is represented as a  
30 percentage of the current resources allocated to a  
31 function.

1        When allocating from the reserve, a percentage of the  
2        reserve is maintained for future allocations. The portion  
3        of the reserve maintained for future allocations can be  
4        established as a fixed percentage or a sliding percentage  
5        that varies as a function of the activity levels of all the  
6        functions, for example, a smaller percentage when the  
7        functions are more active and a larger percentage when the  
8        functions are less active. The resources in the reserve  
9        will slowly decline to zero as the system reaches its  
10       saturation capacity. When the reserve is empty, only  
11       decreases in the activity levels of functions can trigger  
12       reallocations.

13       Resources are proportionally allocated from the  
14       reserve based on the projected increased activity levels of  
15       the functions. In addition, functions having the greatest  
16       projected increase are considered first. For example, if  
17       function A's activity level is projected to increase 30%,  
18       function B's 20%, function C's 10%, and 50 units of the  
19       reserve are eligible to allocate, then function A should  
20       maximally get 25 units ( $1/2 * 50$  units), function B should  
21       maximally get 17 units ( $1/3 * 50$ ), and function C should  
22       maximally get 8 units ( $1/6 * 50$ ).

23       After the maximum allocations have been established,  
24       the allocations and assignment priority are evaluated  
25       against the actual predicted unit increase. For example,  
26       if function A is projected to have a 30% increase but  
27       currently has 10 resource units, then an allocation of 3  
28       units rather than 25 would suffice. Thus, function A is  
29       allocated the 3 units and the difference is left in the  
30       reserve and eligible for allocation. This allows  
31       allocation to function B to be as much as needed.

1       The proportional allocation is then redone as between  
2 functions B and C using the eligible 47 units ( $50 - 3$ ).  
3 Since function B's projected increase was 20% and function  
4 C's projected increase was 10%, function B may be  
5 proportionally allocated 32 units ( $\sim(2/3 * 47)$ ), and  
6 function C may be allocated 15 units ( $47 - 32$ ). However,  
7 if function B is currently allocated 400 units, then a 20%  
8 projected increase implies that an additional 80 units may  
9 be required ( $20\% * 400$ ). However our rule allows only 32  
10 to be allocated. We note however that more were actually  
11 requested and that as many as an additional (80-32) or 48  
12 could be used. If C is currently allocated 10 units, then  
13 a projected 10% increase implies the allocation of 1 unit  
14 ( $10\% * 10$ ). This means that now ( $47-32-1$ ) or 14 units  
15 remain.

16       If any of the functions were allocated fewer resources  
17 than called for by the projection of increased activity  
18 level (for example, function B) and not all of the eligible  
19 resources were allocated in the first pass, then the  
20 process described above is repeated for those functions  
21 having projected activity levels that exceed the resources  
22 allocated in the first pass. In the present example only  
23 function B has an outstanding request for 48 units, and 14  
24 units remain eligible for allocation. Thus, the 14 units  
25 are allocated to function B, leaving a deficiency of 34  
26 units ( $80 - (32 + 14)$ ).

27       It will be appreciated that the inability to allocate  
28 sufficient resources to meet the projected need may or may  
29 not impact performance levels of the functions. If the  
30 projected increased activity levels are not reached, there  
31 is no impact on performance levels of the functions. If  
32 the projections were accurate and a function was not

1 allocated all the resources necessary to meet the  
2 projections, then the performance level of the functions  
3 may be impaired.

4 After the selected resources of the PLD have been  
5 reconfigured in accordance with the selected allocation of  
6 resources, the load monitor continues to sample activity.

7 At step 224, the number of resources maintained in the  
8 reserve is reduced by the amount determined at step 222.  
9 For example, the number of channels that implement a  
10 particular protocol is reduced.

11 At step 228, the allocation of resources removed from  
12 the reserve amongst one or more functions is determined  
13 based on the projected, proportional activity levels of the  
14 functions. The projected activity levels and proportional  
15 allocation are determined as described above in association  
16 with FIG. 3A. Parameters that describe the resources to  
17 allocate to the different functions are then returned to  
18 the process of FIG. 2.

19  
20 FIG. 4 is a flowchart of a process for controlling  
21 reconfiguration of PLDs in response to control parameters  
22 from the load monitor.

23 At step 252, one or more configuration bitstreams are  
24 selected in response to the signals from the load monitor.  
25 The particular configuration bitstreams that are selected  
26 will depend on the PLD resources to be reconfigured and the  
27 selected functions.

28 At step 254, the PLD resources that are to be  
29 reconfigured are selected.

30 The selected configuration bitstreams are then  
31 downloaded to the PLDs at step 256.

1       At step 258, the reconfigured PLDs are restarted with  
2 the reprogrammed functions.

3       Finally, at step 260 the load monitor is signaled that  
4 the reconfiguration is complete.

5       The present invention is believed to be applicable to  
6 a variety of systems having a plurality of PLD-implemented  
7 functions. Other aspects and embodiments of the present  
8 invention will be apparent to those skilled in the art from  
9 consideration of the specification and practice of the  
10 invention disclosed herein. It is intended that the  
11 specification and illustrated embodiments be considered as  
12 examples only, with a true scope and spirit of the  
13 invention being indicated by the following claims.

14